

# QUANTIFYING EVOLUTIONARY CONSTRAINTS ON B CELL AFFINITY MATURATION

CONNOR O. MCCOY, TREVOR BEDFORD, VLADIMIR N. MININ, PHILIP BRADLEY,  
HARLAN ROBINS, AND FREDERICK A. MATSEN IV

**ABSTRACT.** The antibody repertoire of each individual is continuously updated by the evolutionary process of B cell receptor mutation and selection. It has recently become possible to gain detailed information concerning this process through high-throughput sequencing. Here, we develop modern statistical molecular evolution methods for the analysis of B cell sequence data, and then apply them to a very deep short-read data set of B cell receptors. We find that the substitution process is conserved across individuals but varies significantly across gene segments. We investigate selection on B cell receptors using a novel method that side-steps the difficulties encountered by previous work in differentiating between selection and motif-driven mutation; this is done through stochastic mapping and empirical Bayes estimators that compare the evolution of in-frame and out-of-frame rearrangements. We use this new method to derive a per-residue map of selection, which provides a more nuanced view of the constraints on framework and variable regions.

## KEYWORDS

B cell, antibody, immunoglobulin, affinity maturation, molecular evolution, natural selection

## INTRODUCTION

Antibodies encoded by somatically modified human B cell receptor (BCR) genes bind a vast array of antigens, initiating an immune response or directly neutralizing their target. This diversity is made possible by the processes of *VDJ recombination*, in which random joining of V-, D-, and J-genes generates an initial combinatorial diversity of BCR sequences, and *affinity maturation*, which further modifies these sequences. The affinity maturation process, in which antibodies increase binding affinity for their cognate antigens, is essential to mounting a precise humoral immune response. Affinity maturation proceeds via a nucleotide substitution process that combines Darwinian mutation and selection processes. Mutational diversity is generated by *somatic hypermutation* (SHM), in which a targeted molecular mechanism mutates the BCR sequence. This diversity is then passed through a selective sieve in which B cells that bind well to antigen are stimulated to divide, while those that do not bind well or bind to self are marked for destruction. The combination of VDJ recombination and affinity maturation enables B cells to respond to an almost limitless diversity of antigens. Understanding the substitution process and selective forces shaping the diversity of the memory B cell repertoire thus has implications for disease prophylaxis and treatment.

Date: May 11, 2015.

It has recently become possible to gain detailed information about the B cell repertoire using high-throughput sequencing [1–5]. Recent reviews have highlighted the need for new computational tools that make use of BCR sequence data to bring new insight, including the need for reproducible computational pipelines [6–9]. Rigorous analysis of the B cell repertoire will require statistical analysis of how evolutionary processes define affinity maturation. Statistical nucleotide molecular evolution models are often described in terms of three interrelated processes: mutation, the process generating diversity, selection, the process determining survival or loss of mutations, and substitution, the observed process of evolution that follows from the first two processes. One major vein of research has focused on how nucleotide mutation rates depend on the identity of surrounding nucleotides (reviewed in [10]; see also [11, 12]), but little has been done concerning other aspects of the process, such as how the substitution process differs between gene segments.

Along with mutation, selection due to competition for antigen binding forms the other key part of the affinity maturation process. Inference of selective pressures in this context is complicated by nucleotide context-dependent mutation, leading some authors to proclaim that such selection inference is not possible [13]. Indeed, if one does not correct for context-dependent mutation bias, interactions between those motifs and the genetic code can lead to false positive identification of selective pressure. Previous work has developed methodology to analyze selection on sequence tracts in this context (reviewed in Discussion), but no methods have yet achieved the goal of statistical per-residue selection estimates. This has, however, been recently identified as an important goal [11]. Such selection estimates could be used to better direct generation of synthetic libraries of antibodies for high throughput screening. Another application would be to the engineering of antibody Fc regions with specific properties, such as for bispecific monoclonal antibodies or antibody-derived fragments, while preserving overall stability.

The ensemble of germline V, D, and J genes that rearrange to encode antibodies (equivalently: immunoglobulins) are divided into nested sets. They can first be identified by their *locus*: IGH, denoting the heavy chain, IGK, denoting the kappa light chain, or IGL, denoting the lambda light chain. Our dataset contains solely the IGH locus, so we will frequently omit the locus prefix for simplicity. Genes within a locus can be first subdivided by their *segment*, which is whether they are a V, D, or J gene. IGHV genes are further divided into *subgroups* which share at least 75% nucleotide identity. Genes also have polymorphisms that are grouped into *alleles*, which represent polymorphisms of the gene between individuals [14].

VDJ recombination does not always produce a functional antibody, such as when the V and J segments are not in the same reading frame after recombination (an *out-of-frame* rearrangement) or when the receptor sequence contains a premature stop codon. However, each B cell carries two copies of the IGH locus, with one on each chromosome. If the rearrangement on the first locus fails to produce a viable antibody, the second locus will rearrange; if this second rearrangement is successful, the antibody encoded by the second rearrangement will be produced by the cell [15]. If this second rearrangement does not produce a viable antibody the cell dies.

When assaying the BCR repertoire through sequencing, some of the sequences will be from cells for which the first rearrangement was successful, while others will be from cells with one productive and one out-of-frame rearrangement. Although the out-of-frame rearrangements from the second type of cell do not produce viable antibodies, their DNA gets sequenced along with the productive rearrangements. Since SHM rarely introduces insertions or deletions (we observe whole codon insertion deletion events in between 0.013% to 0.014% of memory sequences within templated segments), it is appropriate to assume that observed frame shifts in sequences are dominated by out-of-frame rearrangement events. However, because they are not expressed, but rather are carried along in cells with a separate functional rearrangement, they have no selective constraints. For this reason, we use sequences from out-of-frame rearrangements as a proxy for the neutral mutation process in affinity maturation.

In this paper, we develop modern statistical molecular evolution methods for the analysis of high-throughput B cell sequence data, and then apply them to a very deep short-read data set of B cell receptors. Specifically, we first apply model selection criteria to identify patterns in the single-nucleotide substitution process that occurs during affinity maturation and find that they are similar across individuals but vary significantly across gene segments. Next, we investigate how substitution processes vary between V genes and find that the primary source of variation is whether or not a sequence produces a functional receptor. Finally, we develop the first statistical methodology and corresponding software for comprehensive per-residue selection estimates for B cell receptors. We leverage out-of-frame rearrangements carried along in B cells with a productively rearranged receptor on the second chromosome to estimate evolutionary rates under neutrality, thus avoiding difficulties encountered by previous work in differentiating between selection and motif-driven mutation. A key part of our method is our extension of the “counting renaissance” method for selection inference [16] for non-constant sequencing coverage and a star-tree phylogeny. Using this modified method, we are able to efficiently derive a per-residue map of selection on more than 15 million B cell receptor sequences; we find that selection is dominated by negative selection with patterns that are consistent among individuals in our study.

## RESULTS

**Substitution model inference and testing.** We evaluated the fit of nested models with varying complexity, ranging from a simple model with shared branch lengths and substitution processes for the three independent segments of the BCR, to a complex model with completely separate substitution processes and branch lengths for each segment (Tab. 1a). For the underlying nucleotide substitution model, we fit a general time-reversible (GTR) nucleotide model [17] with instantaneous rate matrix  $Q$  to subsets of the data, using 20,000 unique sequences from each individual. The choice of a stationary and reversible model, rather than a more general model, was based on the similarity of base frequencies between the germline and observed sequences (Tab. S3). We modeled substitution rate heterogeneity across sites using a four-category discretized Gamma distribution [18] with fixed mean 1.0.

We find that the best performing model (denoted  $t_r Q_i \Gamma_i$ , Tab. 1b) is one in

TABLE 1. Models and model testing results.

(A) The models of molecular evolution evaluated, including the number of free parameters introduced in parentheses.

name	branch length	GTR transition matrix	across-site rate variation (discrete Gamma)	total parameters
$t_i Q_i \Gamma_i$	One branch length per <i>segment</i> per sequence ( $n \times 3$ )	One matrix per segment ( $8 \times 3$ )	One distribution per segment (3)	$3n + 27$
$t_r Q_i \Gamma_i$	One branch length per sequence ( $n$ ) + relative rate between segments (2)	One matrix per segment ( $8 \times 3$ )	One distribution per segment (3)	$n + 29$
$t_r Q_i \Gamma_s$	One branch length per sequence ( $n$ ) + relative rate between segments (2)	One matrix per segment ( $8 \times 3$ )	One shared distribution (1)	$n + 27$
$t_r Q_s \Gamma_s$	One branch length per sequence ( $n$ ) + relative rate between segments (2)	One shared matrix (8)	One shared distribution (1)	$n + 11$

(B) Models show identical ranking across individuals. Columns include the log-likelihood (LogL), number of degrees of freedom (df), Akaike Information Criterion (AIC), and difference of AIC from the top model ( $\Delta AIC$ ).

	model	LogL	df	AIC	$\Delta AIC$
A	$t_r Q_i \Gamma_i$	-687,582	20,029	1,415,222	0
	$t_r Q_i \Gamma_s$	-687,980	20,027	1,416,014	793
	$t_r Q_s \Gamma_s$	-700,818	20,009	1,441,654	26,433
	$t_i Q_i \Gamma_i$	-662,417	60,027	1,444,888	29,666
B	$t_r Q_i \Gamma_i$	-507,980	20,029	1,056,017	0
	$t_r Q_i \Gamma_s$	-508,229	20,027	1,056,512	494
	$t_r Q_s \Gamma_s$	-517,320	20,009	1,074,658	18,641
	$t_i Q_i \Gamma_i$	-482,963	60,027	1,085,979	29,962
C	$t_r Q_i \Gamma_i$	-563,181	20,029	1,166,420	0
	$t_r Q_i \Gamma_s$	-563,291	20,027	1,166,637	217
	$t_r Q_s \Gamma_s$	-572,530	20,009	1,185,078	18,659
	$t_i Q_i \Gamma_i$	-539,018	60,027	1,198,090	31,671

which the branch length separating a sequenced BCR from its germline counterpart is estimated independently for each observed sequence, but that V, D and J regions differ systematically in their relative amounts of sequence change (denoted  $t_r$ ). Additionally, this model uses separate GTR transition matrices for V, D and J regions (denoted  $Q_i$ ) and uses separate distributions for across-site rate variation for V, D and J regions (denoted  $\Gamma_i$ ). Looking across models, both the Akaike Information Criterion (AIC) [19] (Tab. 1b) and the Bayesian Information Criterion [20] (data not shown) identified the same rank order of support; this ordering was also identical for each of the three individuals. Other than the  $t_i Q_i \Gamma_i$  model, in which branch length is estimated independently across gene segments, models are ranked in terms of decreasing complexity. The finding that a complex model fits better than simpler models is likely aided by the large volume of sequence data available.




FIGURE 1. Distribution of maximum likelihood branch lengths estimated under the  $t_r Q_i \Gamma_i$  model. Branch lengths are measured in terms of substitutions per site, and rates given for the D and J segments are relative to a fixed rate of 1 for the V segment.

Next, we fit the best-scoring model ( $t_r Q_i \Gamma_i$ ) to the full data set for each individual. The median distance to germline was 0.063, 0.030, and 0.039 substitutions per site for individuals A, B, and C, respectively. The distribution of branch lengths appears nearly exponential for individuals B and C, with many sequences close to germline and few distant from germline sequences (Fig. 1). Individual A displayed a higher substitution load and a non-zero mode. Despite these differences in evolutionary distance, the relative rate of substitution between the V, D, and J segments for each individual was very similar. We note that the sorting procedure used to separate memory from naïve B cells provided memory cells at approximately 97% purity, so these divergence estimates may be conservative due to low levels of contamination from the naïve repertoire.

Coefficients from the GTR models for the same gene segment across individuals were quite similar to one another, while models for different gene segments within an individual showed striking differences (Fig. S1, S2). However, overall correlations of GTR parameters between individuals were very high, yielding correlation coefficients between  $\rho = 0.988$  and  $\rho = 0.994$ . We observe an enrichment of transitions relative to transversions in all segments, as previously described [21].

Next we compared the evolutionary process between various groupings of sequences to learn what determines the characteristics of this evolutionary process. We focused on the V gene segment, as it had the most coverage in our dataset, and partitioned the sequences by whether they were in-frame, then by individual, and then by gene subgroup. We fit the  $t_r Q_i \Gamma_i$  model to 1000 sequences from each set of the partition and calculated the transition probability matrix ( $P$ ) associated with the median branch length across all sequences given an equiprobable starting state. These matrices were then analyzed with a variant of compositional principal components analysis [22] (see Materials and Methods). We find that substitution models are influenced by in- versus out-of-frame sequence status, find no evidence for models clustering by individual, and see some limited evidence for clustering by gene subgroup (Fig. 2). The Euclidean distance between these transformed discrete probability distributions and the Hamming distance




FIGURE 2. First (x-axis) and second (y-axis) principal components from PCA performed on centered log-transformed median-time transition matrices for V gene segments. Points plotted in a random order, with 22 outliers removed for clarity.

between germline V genes showed significant, but moderate, correlation (Spearman's  $\rho = 0.20$ ,  $p < 10^{-15}$ ; Fig. S3).

**Natural selection.** The primary challenge for B cell receptor selection inference is that nucleotide context is known to have a very strong impact on mutation rates (reviewed in [21]). These context-specific mutations combined with the structure of the genetic code can result in extreme  $dN/dS$  ratios using the classical definition that are not attributable to selection. To address this problem, we infer the selection coefficient  $\omega$  using a nonsynonymous-synonymous ratio which controls for background mutation rate via out-of-frame sequences (3). We continue the tradition of calling the selection coefficient  $\omega$  in this context, even though it is a slightly different definition than previously used.

Applying this method to our data set results in the first per-site and per-gene maps quantifying selection in the B cell repertoire [23, 24]. Sites were classified as negatively or positively selected based on whether the 95% Bayesian credible interval (BCI) excludes 1.0: sites for which the lower endpoint of the  $\omega$  BCI is greater than 1.0 are classified as being under positive selection, while sites for which the upper endpoint of their  $\omega$  BCI is less than 1.0 are classified as being under negative selection. We employ site numbering according to the IMGT unique numbering for the V domain [25].

IGHV3-23\*01 is the most frequent V gene/allele combination in our dataset, and it displays patterns that are consistent with the other genes. Specifically, we see significant variation in the synonymous substitution rate (right panels, Fig. 3a) even in out-of-frame sequences, which is presumably due to motif-driven mutation. Thus, if we had directly applied traditional means of estimating selection by comparing the rate of nonsynonymous and synonymous substitutions, we would have falsely identified sites as being under strong selection. In contrast the selection inferences made using out-of-frame sequences stay much closer to neutral (Fig. 3b).

We note extensive negative selection in the residues immediately preceding the CDR3 (Fig. 4). The amino acid profile for these sites shows a distinct preference

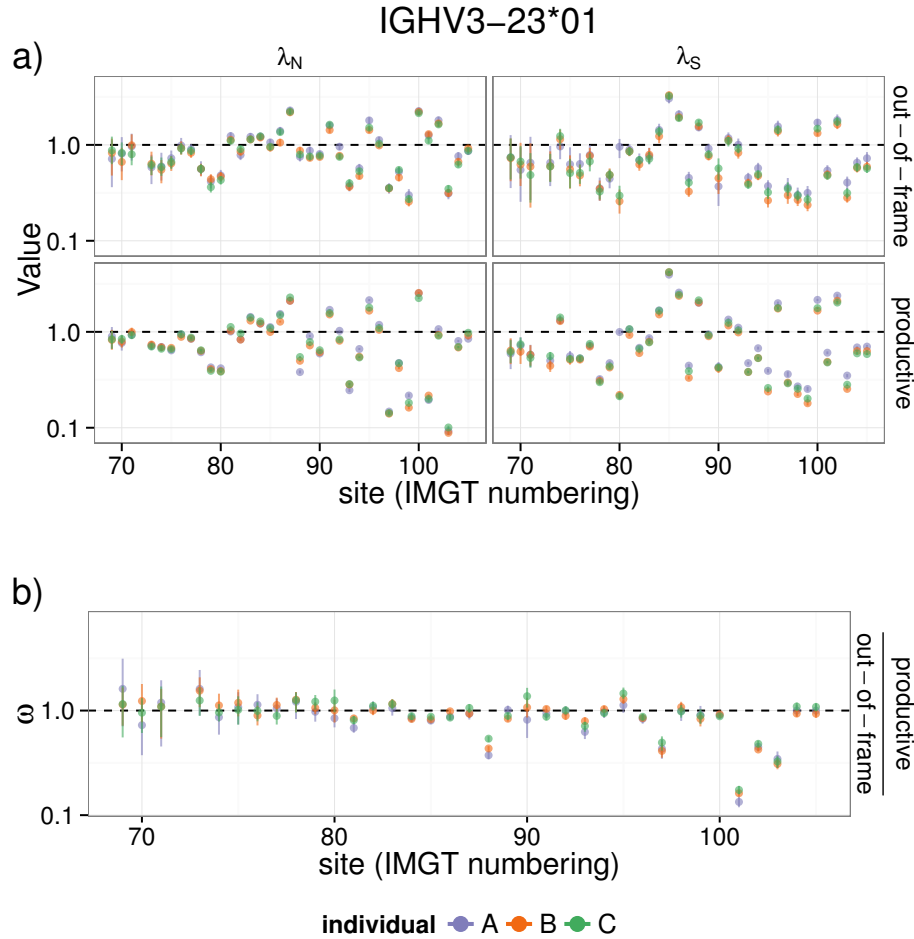


FIGURE 3. a) Comparison of nonsynonymous ( $\lambda_N$ ) and synonymous ( $\lambda_S$ ) rates in productive and out-of-frame sequences. b)  $\omega$  estimates using unproductive rearrangements as a proxy for the neutral process. Both panels use data from IGHV3-23\*01, the most frequent V gene/allele combination.

for a tyrosine or rarely a phenylalanine two residues before the start of the CDR3 at site 102 (Fig. S4). It shows a preference for a tyrosine or more rarely a phenylalanine or a histidine in the residue just before the start of the CDR3 at site 103. These aromatic positions likely play important structural roles in the antibody complex: site 102 is buried in the core of the heavy chain and makes extensive van der Waals interactions as well as a sidechain-backbone hydrogen bond, while site 103 forms part of the interface between the heavy and light chains (see further description of structural results below).

Overall we see extensive selection in our sequenced region (Fig. 5). The mean  $\omega$  estimate across sites with at least 100 productive and out-of-frame sequences




FIGURE 4. Site-specific estimates of the selection coefficient  $\omega$ . Violin plots show distribution of median  $\omega$  estimates across V genes at each site. Bar plots show count of V genes classified as undergoing negative, neutral, or positive selection. Only sites with at least 100 productive and out-of-frame observed sequences aligned were considered. The sites with IMGT numbers less than or equal to 104 are traditionally designated “framework.”



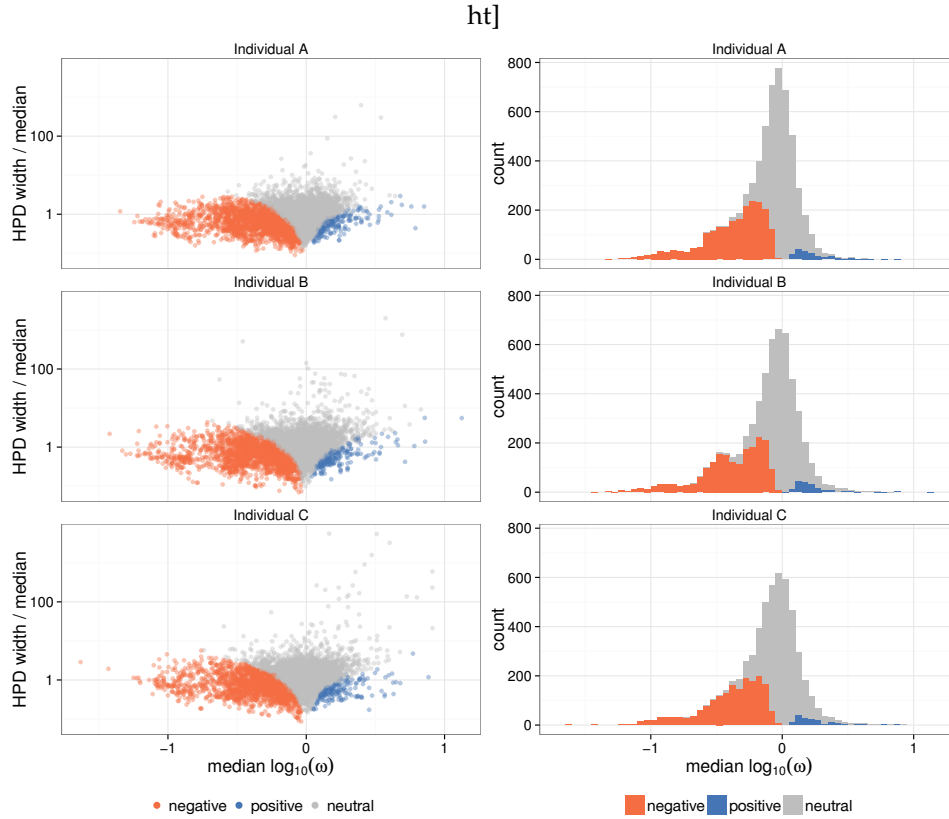


FIGURE 5. Site-specific selection estimates partitioned by individual and gene. Sites classified as negatively selected or positively selected based on whether the 95% Bayesian credible interval excludes 1 and in what direction. Left: comparison of  $\omega$  estimate and relative width of BCI region. Right: distribution of site-specific selection estimates.

aligned was 0.907. 65.6% of sites had a median  $\omega < 1$  with a wide distribution of median  $\omega$  values and confidence interval widths. However, many of them were observed to be positively, negatively, and neutrally evolving with narrow confidence intervals (Fig. 5, left column). 30.6% of sites were confidently classified as being under negative selection (Fig. 5, right column).

Because amino acids interior to the protein could be important for protein stability compared to exposed ones, we hypothesized that residues under negative selection would be more internal to the antibody protein than those under neutral or positive selection, and that the inverse would be true for residues under positive selection. To test this, we mapped our  $\omega$  estimates onto antibody structures (Fig. 6) and calculated the exposure of each amino acid position in the structure




FIGURE 6. An IGHV3-23\*01 (the most frequent V gene/allele combination) heavy chain antibody in complex with IL-17A (PDB ID 2VXS; [28]), with sites colored by  $\omega$  classification in each of the three individuals sampled. The bound antigen is shown in pink (top), and the light chain in green (right). The heavy chain structure is shown as a thin gray line at sites which could not be classified due to insufficient coverage. When there is sufficient coverage it is shown as a cartoon (thick lines or arrows representing beta sheets) which is colored gray at neutral sites, red at negatively selected sites, and blue at positively selected sites.

using the solvent-accessible surface area (SASA) using ROSETTA3 [26]. The normalized SASA was well correlated with the classification of each site: sites classified as being under positive selection were most exposed in the protein structure, followed by neutral sites, then negatively selected sites (Fig. S7). Differences in surface accessibility were significant between the three groups, with p-values ranging from  $< 0.002$  for the comparison of positive vs. neutral sites to  $< 10^{-15}$  for the comparison of negative vs. neutral sites (Wilcoxon rank-sum test [27]).

Despite the three individuals surveyed here presumably having quite different immune histories, we observe remarkable consistency in substitution and selection within the memory B cell repertoire. Indeed, we see a very strong correlation of median selection estimates between individuals (Fig. S5), with between-individual coefficients of determination  $R^2$  of between 0.628 and 0.687 for site-specific  $\omega$  values.

## DISCUSSION

We find different patterns of substitution across the V, D and J regions which is consistent among individuals (Fig. S1) even though those individuals have differing levels of substitution (Fig. 1). We find that the dominant factor determining the V segment substitution process is whether it is out-of-frame or productive,

with the gene identity being a contributing factor. The pattern of selective pressure is consistent across individuals, and shows especially strong pressure near the boundary between the V gene and the CDR3. Selection estimates for BCRs are still high, with average  $\omega$  of  $\approx 0.9$ , compared to common examples of Darwinian evolution, such as seen in *Drosophila* [29] and mammals [30], where most genes show  $\omega$  less than 0.1. However, we note that although our estimates of  $\omega$  are comparable to more traditional estimates, we calculate  $\omega$  slightly differently, using out-of-frame sequences as a control for motif-driven evolution. Finally, the patterns of selective pressure we observed correlated with levels of surface exposure in published antibody structures: highly conserved sites were more frequently found internally, while residues we classified as positively selected were more exposed.

We note that our analyses are based on data from only three individuals. It is possible that including more individuals would reveal variation in the mutation process. However, we note that these three unrelated individuals had an extraordinary level of agreement, which cannot be explained by sequencing error.

**Substitution process.** We closely examine the substitution and selection processes in a context-independent manner, not to make a full description of this clearly context-dependent process, but rather to provide a solid framework for future study and to enable downstream comparative analyses (Figure 2). Our model selection shows that the best-fitting model allows for a single branch length per sequence, a global multiplier for per-segment differences, a per-segment substitution model, and a per-segment rate variation model across sites (Tab. 1b). These between-segment differences are certainly due in part to base composition, which also differs significantly between segments and is similar between individuals (Table S3). Another contributing factor is probably similarity of local nucleotide context between the genes of a given segment compared to between segments; these nucleotide contexts are known to impact AID-induced somatic hypermutation (reviewed in [21]). We also note that the entirety of the D segment lies within the CDR3 region, and is thus more likely to directly contact an epitope; not surprisingly, we observe higher substitution rates within that segment. By analyzing distances between GTR substitution rate matrices, we find that the most important difference between them is determined by whether they are productive or non-productive (Fig. 2), which is presumably due to the impact of natural selection. We also find a significant correlation between sequence identity and substitution matrix (compare [31]). In a related though distinct vein, [32] develop an amino acid substitution model for BCR sequences, which analogously aggregates information across positions.

**Selection process.** The role of selection in B cell receptor development has stimulated continuous interest since the pioneering 1985 paper of Clarke and colleagues [33], however methods for the analysis of antigen selection have developed in parallel to related work in the population genetics and molecular evolution community. Work on the selection process for BCRs has focused on aggregate statistics to infer selection for entire sequences or sequence tracts, and there has been a lively debate about the relative merits of these tests [34–38]. Recent work has offered methods that evaluate selection on a per-sequence basis [38]. There have also been

efforts to infer selection based on lineage shape [39–44], which has been a common approach in macroevolutionary studies (reviewed in [45]) and more recently in population genetics [46–49].

In this work, we develop the first means of inferring per-residue selection using high-throughput sequence data with non-uniform coverage. Our method sidesteps the difficulties encountered by previous work in differentiating between selection and motif-driven mutation in B cell receptors [11, 13, 34–38] by developing statistical means to compare in-frame and out-of-frame rearrangements. An alternate means of estimating selection was recently developed by [12] in which a regression model incorporating a detailed model of motif preferences was used to infer selection coefficients for the framework and CDR regions in aggregate. In contrast to this previous work on B cell selection, our methods provide a *per-residue* selection map for a contiguous stretch of BCR sequence.

We use out-of-frame rearrangements as our selection-free control population. These sequences do not create functional IGH proteins, but may be carried in heterozygous B cells which do have a productively rearranged IGH allele. Thus they undergo SHM, but any selection occurs on the level of the productively rearranged allele, not on the residues in the unproductive allele. We observe a similarly high proportion of germline-identical sequences for in-frame and out-of-frame subsets in naïve cells (Table S2); differences from germline derive in part from sequencing and other platform errors that do not depend on frame. For memory cells, we see extensive action of somatic hypermutation, but with a higher proportion of germline-identical out-of-frame sequences than in-frame (Table S2). We interpret these additional mutations for in-frame memory sequences as following from the process of affinity maturation for a specific antigen. We acknowledge that some out-of-frame sequences could still feel the impact of selection, which would occur if the sequences accrue frameshift mutations in the process of affinity maturation. However, it is thought that SHM is primarily a process of point mutation [21], and indeed, we observe whole codon indels in only 0.013%–0.014% of memory sequences within templated segments. Still, if a weaker version of selection was occurring on the out-of-frame sequences compared to the productive ones then this would simply make our estimates of selection conservative, pulling estimates of  $\omega$  closer to 1, and yet our selection estimates are confidently classified as non-neutral for a substantial fraction of sites (Fig. 5).

In applying our methodology to IGHV sequences, we gain a high resolution per-gene map of selective forces on B cell receptors for part of the V gene. This part is primarily in the framework region, which is thought to be under substantial evolutionary constraint to preserve structure. Indeed, we see an pattern of quite strong negative selection in the region around the beginning of the CDR3 (Figure 4), agreeing with recent work that found strong negative selection in one site near the beginning of the CDR3 [11]. However, other sites in this section of framework have substantially relaxed selection (Figure 4). These results thus provide a more nuanced view into the constraints on B cell receptor sequences rather than the traditional framework/variable designations, as also noted by [11].

This work points the way towards future directions. In this work we assumed that the size of individual lineages is small compared to the size of the overall repertoire, and thus that lineage structure could be ignored for the purpose of evolutionary model analysis. Ideally we would reconstruct lineages and then do




FIGURE 7. B cell receptor schematic showing variable (V), diversity (D) and joining (J) gene segments as well as framework (FW) and complementarity-determining regions (CDRs). In VDJ recombination, individual V, D, and J gene segments are randomly selected and joined together via a process that deletes some randomly distributed number of nucleotides on their boundaries then joins them together with random “non templated insertions” (N). The specificity of an antibody is primarily determined by the region defined by the heavy chain recombination site, referred to as the third complementarity determining region (CDR3). The sequence data for this study started in the fourth framework (FW) region and continued into the third. Amplification was via a forward primer in the FW2 region and a reverse primer in the FW4 region.

evolutionary analysis on the tree corresponding to each lineage. However, reconstructing groups of sequences forming a lineage is a challenging prospect on its own, to say nothing of doing phylogenetics on sequences in the presence of strong context-specific mutation-selection patterns, and have left out incorporating those aspects until we have first developed the necessary methods. We have recently developed an HMM framework to analyze VDJ rearrangements [50], and are currently developing and validating ways to use this framework for likelihood-based (as opposed to procedure-based [51,52]) lineage group inference.

## MATERIALS AND METHODS

**Data set.** The complete description of the experiment will be published separately (manuscript in preparation). However, here we give a brief overview of the data in order to facilitate understanding of our analysis and to emphasize that the experimental design has a number of features that greatly reduce errors in sequencing and quantification. 400cc of blood was drawn from three healthy volunteers under IRB protocol at the Fred Hutchinson Cancer Research Center. CD19<sup>+</sup> cells were obtained by bead purification then flow sorted to isolate over 10 million naïve (CD19<sup>+</sup>CD27<sup>-</sup>IgD<sup>+</sup>IgM<sup>+</sup>) and over 10 million memory (CD19<sup>+</sup>CD27<sup>+</sup>) B cells, with greater than 97% purity. Genomic DNA was extracted and the ImmunoSeq assay described in [3] was performed on the six samples at Adaptive Biotechnologies in Seattle, WA.

The experiments and preprocessing were carefully designed to give an accurate quantification of error-corrected observed sequences. To mitigate preferential amplification of some V/J pairs through primer bias, the PCR amplification was performed using primers optimized via a large collection of synthetic templates [53].

To reduce sequencing errors and provide accurate quantification, each sample was divided amongst the wells on two 96 well plates and bar-coded individually. These templates were then amplified and “over-sequenced” (Table S1), so that an average of almost 6 reads were present for each template. Following [54], reads matching the same template were collapsed into a consensus sequence with reduced sequencing error. Here, we grouped reads from within a well into consensus sequences by joining reads with Hamming distance less than or equal to two, and inferred the consensus sequence in each group using parsimony. Groups with only one member were discarded. This procedure protects against collapsing distinct sequences, as the probability that nearly identical distinct sequences co-occur exclusively in the same wells is small. We acknowledge this procedure may eliminate low frequency variants, but we prioritized accuracy over sensitivity towards these variants; note that despite this conservative analysis pipeline we observed substantial signal in the data.

Deep sequencing these B cell populations resulted in 15,023,951 (Tab. S1) unique 130bp observed sequences after pre-processing that spanned the third heavy chain complementarity determining region (CDR3) region (Fig. 7). The full data set will be made public upon publication of the manuscript describing the experiment.

**Alignment and germline assignment.** Each sequence was first aligned to each V gene using the Smith-Waterman algorithm with an affine gap penalty [55]. The 3' portion of the sequence not included in the best V gene alignment was next aligned to all D and J genes available from the IMGT database [14]. The best scoring V, D, and J alignment for each sequence was taken to be the germline alignment, and the corresponding germline sequence was taken to be the ancestral sequence for that observed sequence; in the case of ties, one germline sequence was chosen randomly among those alleles present at abundance  $\geq 10\%$ . Sequences were classified as productive or out-of-frame based on whether the V and J segments were in the same frame; all sequences with stop codons were removed, as these sequences could result from either an unproductive rearrangement event or inactivation due to a lethal mutation. The 18 V gene polymorphisms present at the highest frequency in the naïve populations of the individuals surveyed which were not represented in the IMGT database were added to the list of candidates for alignment. In contrast to naïve sequences which have no mutations across almost all sites, the alleles we added to the germline collection were all present at greater than 30% for the IGHV gene in question.

**Substitution models, fitting and analysis.** The setting of B cell affinity maturation is substantially different than that typically encountered in molecular evolution studies, and hence there are some differences between our model fitting procedure compared to common practice. For B cell receptors outside of nontemplated insertions, the root state is the V, D, and J genes encoded in the germline from which a sequenced BCR derives. Thus, we analyze substitutions that have occurred in evolution away from the germline-encoded segments of observed BCR sequences, ignoring sites comprising nontemplated insertions. The CDR3 region of an antibody is generally sufficient to uniquely identify its specificity [56]. Although there are certainly some clones in our data set that derive from a single rearrangement event but differ due to somatic hypermutation, the probability that a given pair of distinct sequences derives from a single common ancestor is small:

targeted searches for clonally related antibodies during infection have identified them at 0.003% to 0.5% [57]. It is a substantial challenge to statistically infer which sequences sit together in a clonal lineage and then to perform phylogenetic analysis on such a large data set (see work by [12, 58, 59]) and performing this analysis incorrectly could bias our results. Additionally, we encountered significant computational barriers analyzing the volume of sequences available, and adding phylogenetic structure to our analysis may have made the analysis computationally prohibitive even if we had the lineage structure in hand (we believe this is the largest number of sequences from a single data set analyzed in selection study to date).

For these reasons, our analyses were performed on a set of pairwise alignments, each representing a two taxon tree containing an observed sequence and its best scoring germline sequence according to Smith-Waterman alignment. This is equivalent to using a rooted “star” tree where the root state is known. This assumption allowed us to focus our attention on the selection inference problem.

Substitution models are summarized in Table 1a and described in detail here. We will use  $n$  for the number of observed sequences. Our models are characterized by three components. First, the subscript of  $t$  describes how branch length assignments are allowed to vary across segments of a single sequence. The  $t_i$  model allows branch lengths to vary independently, resulting in  $3n$  parameters. The  $t_r$  model has two global per-segment multipliers to define the branch lengths (see, e.g. Fig. 1) with the V segment rate fixed at 1, resulting in  $n + 2$  parameters. The subscript of  $Q$  describes how rate matrices are fit. The  $Q_i$  model allows an independent global GTR rate matrix for each segment, with a total of 24 parameters. The  $Q_r$  model just has one GTR rate matrix overall, with 8 parameters. The subscript of  $\Gamma$  denotes how across-site substitution rate variation is modeled in terms of a four category discrete gamma distribution [18]. The  $\Gamma_i$  model allows an independent rates across sites parameter for each sequence, with 3 parameters. The  $\Gamma_s$  has a global rates across sites parameter, with 1 parameter. Given these choices concerning how the data was partitioned and parametrized, the standard phylogenetic likelihood function was used as described in the original literature [60–62] and in books (e.g. [63, 64]).

Maximum likelihood values of substitution model parameters and branch lengths were estimated using a combination of Bio++ [65] and BEAGLE [66], with model optimization via the BOBYQA algorithm [67] as implemented in NLOpt [68], and branch length optimization via Brent’s method [69]. Optimization alternated between substitution model parameters and branch lengths until the change in log-likelihood at a given iteration was less than 0.001. Our software to perform this optimization is available from <https://github.com/cmccoy/fit-star>.

For the principal components analysis on substitution matrices, we first obtained the median branch length  $\hat{t}$  across all sequences for all individuals. We then calculated the corresponding transition matrix for each model given equiprobable starting state:  $e^{Q\hat{t}} \text{diag}(0.25)$ . These were then projected onto the first two principal components, adapting suggestions for doing PCA in the simplex [22]. Specifically, each row of these matrices, as a discrete probability distribution, is a point in the simplex. Hence we applied a centered log transformation to each row of this matrix using the `clr` function of the R package `compositions` [70], and followed with standard principal components analysis.

To compare distance between inferred models and sequence distance, we calculated the Hamming distance between all pairs of V genes using the alignment available from the IMGT database [14]. To obtain distances between models, we calculated the Euclidean distance calculated between pairs of the transformed probability vectors.

**Selection analysis.** Because of the large volume of sequences to analyze, we also needed a mechanism to detect selection that could be run on over 15 million sequences, most of which did not share evolutionary history. Classical means of estimating selection by codon model fitting [71, 72] could not be used, even in their most recent and much more efficient incarnation [73]. Instead, we used the renaissance counting approach [16], which we modified to work under varying levels of coverage. A key part of the renaissance counting approach is an empirical Bayes regularization procedure [74]. This procedure uses the entire collection of sites to inform substitution rate estimation at each site individually, effectively sharing data across sites, allowing inference at sites which either display few substitutions or have less sequencing coverage. We note that obtaining precise per-site selection estimates for hundreds of genes requires a large quantity of sequence data like what we have here: the read coverage decrease on the 5' end of the V gene correspondingly increases the width of the error bar (Figure 3, see [23]), resulting in a decrease of power for selection regime classification (Figure 4).

*Bayesian inference of selection on a star-shaped phylogeny.* To determine the site-specific selection pressure for each V gene, we extended the counting renaissance method, described in [16], to accommodate pairwise analyses of a large number of sequences with a known ancestral sequence and non-constant site coverage. The counting renaissance method starts by assuming a separate HKY substitution model [75] for each of the three codon positions and uses Markov chain Monte Carlo (MCMC) to approximate the posterior distribution of model parameters that include substitution rates and phylogenetic tree with branch lengths. Since in our analyses we assumed that query sequences are related by a star-shaped phylogeny, our model parameters included only HKY model parameters and branch lengths leading to all the query sequences. Moreover, we fixed the parameters of the HKY model, along with the relative rates between codon positions, to the maximum likelihood estimates produced using the whole dataset. We note that we could have fit per-codon-position GTR models and used them for stochastic mapping, however such a model would still be substantially misspecified compared to a codon model and thus we decided to follow [16] and use HKY for the mapping. *A priori*, we assumed that branch lengths leading to the query sequence independently follow an exponential distribution with mean 0.1. We performed 20,000 iterations of MCMC, scaling the branch length leading to the observed sequence at each iteration, and sampling every 40 iterations to generate a total of 500 samples. Given each posterior sample of query branch lengths, the counting renaissance method draws a sample of ancestral substitutions conditional on the observed data using a simple per-codon-position nucleotide model; the resulting sampled ancestral sequences are then used to count synonymous and nonsynonymous mutations.



*Sampling codon substitutions.* For each unique read, for each codon position  $l$  and posterior sample  $j$ , counts of synonymous ( $C_{jl}^{(S)}$ ) and nonsynonymous ( $C_{jl}^{(N)}$ ) substitutions at each site were imputed using stochastic mapping as described above.

For  $N$  MCMC iterations based on an alignment of  $L$  codons, the result of this procedure was two  $N \times L$  matrices, each containing the number of synonymous and nonsynonymous events at each codon position in each posterior sample. Counts of each substitution type along with the total branch length for each site were aggregated across sequences from the same gene by element-wise addition. This took about 5 days on an 194 core cluster launched on Amazon Web Services using starcluster (<http://star.mit.edu/cluster/>).

*Empirical Bayes regularization.* The varying length of the CDR3, combined with short observed sequences, leads to quite skewed coverage of sites stratified by gene. We modified the empirical Bayes regularization procedure of the original counting renaissance method [16] to account for varying depth of observation as follows. First, we define a branch length leading to query sequence  $i$  for site  $l$  as

$$t_{il} = \begin{cases} t_i, & \text{if any residues in the observed sequence } i \\ & \text{align to codon position } l \\ 0, & \text{otherwise} \end{cases}$$

We assume that substitution counts for site  $l$  come from a Poisson process with rate  $\lambda_l t_l$ :

$$C_l \sim \text{Poisson}(\lambda_l t_l),$$

where  $t_l = \sum_{i=1}^n t_{il}$ .

As in the original counting renaissance, we assume that the site-specific rates  $\lambda_l$  come from a Gamma distribution with shape  $\alpha$  and rate  $\beta$ :

$$\lambda_l \sim \text{Gamma}(\alpha, \beta).$$

We fix  $\alpha$  and  $\beta$  to their maximum likelihood estimates  $\hat{\alpha}$  and  $\hat{\beta}$  by treating sampled branch lengths and counts as fixed and maximizing the likelihood function

$$(1) \quad \mathcal{L}(\alpha, \beta) = \left( \frac{\beta^\alpha}{\Gamma(\alpha)} \right)^L \prod_l \frac{t_l^{C_l}}{\Gamma(C_l + 1)} \frac{\Gamma(C_l + \alpha)}{(t_l + \beta)^{C_l + \alpha}}.$$

We provide a derivation of this likelihood function in the Supplementary Methods. In contrast to [16], we do not have closed-form solutions for the maximum likelihood or method of moments estimators of  $\alpha$  and  $\beta$  in this slightly more complex setting. However it does not add a substantial computational burden to estimate these parameters numerically via the BOBYQA optimizer [67].

Given  $\hat{\alpha}$  and  $\hat{\beta}$ , we draw rates  $\lambda_l$  from the posterior:

$$(2) \quad \lambda_l \mid C_l \sim \text{Gamma}(C_l + \hat{\alpha}, t_l + \hat{\beta}),$$

derived in the Supplementary Methods.

Estimation of  $\alpha$  and  $\beta$  by maximizing likelihood (1) fails when the sample variance of the observed counts  $C_1 \dots C_L$ , weighted by the site-specific branch length sums,  $t_1 \dots t_L$ , is less than the corresponding weighted sample mean. In these

cases, we assume that the observed counts are drawn from Poisson distributions with site-specific rate  $\lambda t_l$ :

$$C_l \sim \text{Poisson}(\lambda t_l),$$

where here  $\lambda$  is shared across sites, and is estimated from the data by maximizing the likelihood

$$L(\lambda) = \prod_l \frac{(\lambda t_l)^{C_l}}{C_l!} e^{-\lambda t_l}.$$

*Simulations.* To validate this method, we simulated 1000 sequences of 100 codon sites each under the GY94 model and a star-like phylogeny with branch lengths fixed to 0.05 using piBUSS [76]. We varied  $\omega$  over the alignment, with 85 sites having  $\omega = 0.1$ , 5 sites having  $\omega = 1$ , and 10 sites under positive selection -  $\omega = 10$ . We next introduced varying coverage over the alignment: sequences were truncated such that no sequences covered the first 10 codons, only half of the sequences had coverage over the next 40 codons, and all sequences covered the remaining 50 codons (Fig. S6, bottom panel). Estimates of  $\omega$  were more accurate with higher site coverage (Fig. S6, top panel). Of note, as a result of the empirical Bayes regularization, even some sites with no coverage were classified as being under purifying selection. In all other analyses, we only report  $\omega$  estimates for sites covered by at least 100 sequences. Since the starting state is always the germline amino acid, no classifications can be made for sites which are Tryptophan or Methionine in the germline, as all mutations are nonsynonymous for codons encoding those amino acids.

*Site-specific estimates of  $\omega$ .* In [16], the authors arrive at site-specific estimates of  $\omega_l$  by comparing data-conditioned ( $C$ ) rates  $\lambda_l$  of nonsynonymous ( $N$ ) and synonymous ( $S$ ) substitutions, each normalized by an “unconditional rate” ( $U$ ):  $\omega_l^{RC} = \frac{\lambda_l^{(N,C)} / \lambda_l^{(N,U)}}{\lambda_l^{(S,C)} / \lambda_l^{(S,U)}}$ . As SHM is highly context-specific, we chose to use rates inferred from out-of-frame rearrangements in place of the unconditional rates, as these more accurately represent the mutation rates in the absence of selection:

$$(3) \quad \omega_l = \frac{\lambda_l^{(N,P)} / \lambda_l^{(N,O)}}{\lambda_l^{(S,P)} / \lambda_l^{(S,O)}},$$

where  $P$  and  $O$  refer to productive and out-of-frame rearrangements, respectively.

*Implementation.* We used the BEAST [77] implementation of the counting renaissance procedure to sample counts for both synonymous and nonsynonymous substitutions at each site. We extended BEAST version 1.8.0 to generate “unconditional” counts using the germline state as the starting state for simulating along the edge to the query as described above. This process (sampling substitutions for each sequence, then combining counts from sequences mapping to the same IGHV) provides a natural setting for parallelization via the map-reduce model of computation; we used the Apache Spark [78] framework to distribute work across a cluster running on Amazon EC2. Our software to perform this analysis is available from <https://github.com/cmccoy/startreerenaissance>.

**Structural analysis.** For each of the eleven most frequently occurring V genes, we identified the closest structure in the Protein Data Bank (PDB) [79] using BLAST [80]. Structures were visualized using PyMOL [81]. We calculated the normalized solvent-accessible surface area (SASA) for each amino acid position using ROSETTA3 [26] and normalized these values by dividing them by the fully exposed SASA of the given residue type in an extended chain. Wilcoxon rank-sum tests [27] between all pairs of selection classifications (negative, neutral, positive) were used to assess whether the normalized SASA differed between groups. p-values were Bonferroni-corrected [82] to account for multiple comparisons.

The details of our computational methods are available in the *Supplementary Methods*.

#### DATA ACCESSIBILITY

Data is available at <http://adaptivebiotech.com/link/mat2015>. We have made an Amazon Machine Image (AMI) [83] available pre-loaded with our analysis pipeline and some example data. To use it, launch an instance of AMI `ami-ab295b9b` in the `us-west-2` region and log in as user `ubuntu` (no password needed: authentication by public key).

#### FUNDING

This work was supported by the National Institute of Health [R01-AI107034 to V.N.M., R01-GM11324601 to F.A.M., R01-AI103981 to Dr. Julie Overbaugh] and the National Science Foundation [DMS-0856099 to V.N.M.]; C.O.M. and F.A.M. supported in part by the University of Washington eScience Institute through its Seed Grants program in Translational Health Sciences, and by a 2013 new investigator award from the University of Washington Center for AIDS Research (CFAR), an NIH funded program under award number P30AI027757 which is supported by the following NIH Institutes and Centers: NIAID, NCI, NIMH, NIDA, NICHD, NHLBI, NIA, NIGMS, and NIDDK.

#### AUTHORS' CONTRIBUTIONS

CM designed and carried out the computational analysis, developed the statistical methods, and wrote the paper; TB designed the computational analysis, developed the statistical methods, and wrote the paper; VM designed the computational analysis, developed the statistical methods, and wrote the paper; PB performed structural analysis; HR helped define questions and provided the data; FM designed the computational analysis, developed the statistical methods, and wrote the paper.

#### COMPETING INTERESTS

HR owns stock in and consults for Adaptive Biotechnologies. The other authors have no competing interests.

## ACKNOWLEDGEMENTS

The molecular work for this project was done by Paul Lindau in the laboratory of Phil Greenberg, and was supported by a grant from the W. M. Keck Foundation. The authors are grateful to Bryan Howie and William DeWitt III for helpful discussions.

## REFERENCES

- [1] Boyd SD, Marshall EL, Merker JD, Maniar JM, Zhang LN, Sahaf B, et al. Measurement and Clinical Monitoring of Human Lymphocyte Clonality by Massively Parallel V-D-J Pyrosequencing. *Sci Transl Med*. 2009 Dec;1(12):12ra23–12ra23. Available from: <http://dx.doi.org/10.1126/scitranslmed.3000540>.
- [2] Wu YC, Kipling D, Leong HS, Martin V, Ademokun AA, Dunn-Walters DK. High-throughput immunoglobulin repertoire analysis distinguishes between human IgM memory and switched memory B-cell populations. *Blood*. 2010 10 May;116(7):1070–1078. Available from: <http://dx.doi.org/10.1182/blood-2010-03-275859>.
- [3] Larimore K, McCormick MW, Robins HS, Greenberg PD. Shaping of human germline IgH repertoires revealed by deep sequencing. *J Immunol*. 2012 3 Aug;189(6):3221–3230. Available from: <http://dx.doi.org/10.4049/jimmunol.1201303>.
- [4] DeKosky BJ, Ippolito GC, Deschner RP, Lavinder JJ, Wine Y, Rawlings BM, et al. High-throughput sequencing of the paired human immunoglobulin heavy and light chain repertoire. *Nat Biotechnol*. 2013 20 Feb;31(2):166–169. Available from: <http://dx.doi.org/10.1038/nbt.2492>.
- [5] Robins H. Immunosequencing: applications of immune repertoire deep sequencing. *Curr Opin Immunol*. 2013 16 Oct;25(5):646–652. Available from: <http://dx.doi.org/10.1016/j.coi.2013.09.017>.
- [6] Mehr R, Sternberg-Simon M, Michaeli M, Pickman Y. Models and methods for analysis of lymphocyte repertoire generation, development, selection and evolution. *Immunol Lett*. 2012 10 Aug;148(1):11–22. Available from: <http://dx.doi.org/10.1016/j.imlet.2012.08.002>.
- [7] Six A, Mariotti-Ferrandiz ME, Chaara W, Magadan S, Pham HP, Lefranc MP, et al. The Past, Present, and Future of Immune Repertoire Biology - The Rise of Next-Generation Repertoire Analysis. *Front Immunol*. 2013 27 Nov;4:413. Available from: <http://dx.doi.org/10.3389/fimmu.2013.00413>.
- [8] Warren EH, Matsen FA IV, Chou J. High-throughput sequencing of B- and T-lymphocyte antigen receptors in hematology. *Blood*. 2013 8 May;122(1):19–22. Available from: <http://dx.doi.org/10.1182/blood-2013-03-453142>.
- [9] Georgiou G, Ippolito GC, Beausang J, Busse CE, Wardemann H, Quake SR. The promise and challenge of high-throughput sequencing of the antibody repertoire. *Nat Biotechnol*. 2014 19 Jan; Available from: <http://dx.doi.org/10.1038/nbt.2782>.
- [10] Delker RK, Fugmann SD, Papavasiliou FN. A coming-of-age story: activation-induced cytidine deaminase turns 10. *Nat Immunol*. 2009 20 Oct;10(11):1147–1153. Available from: <http://dx.doi.org/10.1038/ni.1799>.
- [11] Yaari G, Vander Heiden JA, Uduman M, Gadala-Maria D, Gupta N, Stern JNH, et al. Models of somatic hypermutation targeting and substitution based on synonymous mutations from high-throughput immunoglobulin sequencing data. *Front Immunol*. 2013 15 Nov;4:358. Available from: <http://dx.doi.org/10.3389/fimmu.2013.00358>.
- [12] Kepler TB, Munshaw S, Wiehe K, Zhang R, Yu JS, Woods CW, et al. Reconstructing a B-Cell Clonal Lineage. II. Mutation, Selection, and Affinity Maturation. *Front Immunol*. 2014 22 Apr;5:170. Available from: <http://dx.doi.org/10.3389/fimmu.2014.00170>.
- [13] Dunn-Walters DK, Spencer J. Strong intrinsic biases towards mutation and conservation of bases in human IgVH genes during somatic hypermutation prevent statistical analysis of antigen selection. *Immunology*. 1998 Nov;95(3):339–345. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/9824495>.
- [14] Lefranc MP, Giudicelli V, Ginestoux C, Jabado-Michaloud J, Folch G, Bellahcene F, et al. IMGT, the international ImMunoGeneTics information system. *Nucleic Acids Res*. 2008 31 Oct;37(Database issue):D1006–12.
- [15] Corcoran AE. Immunoglobulin locus silencing and allelic exclusion. *Semin Immunol*. 2005 Apr;17(2):141–154.
- [16] Lemey P, Minin VN, Bielejec F, Kosakovsky Pond SL, Suchard Ma. A counting renaissance: combining stochastic mapping and empirical Bayes to quickly detect amino acid sites under positive selection. *Bioinformatics*. 2012 Dec;28(24):3248–3256. Available from: <http://dx.doi.org/10.1093/bioinformatics/bts580>.

- [17] Lanave C, Preparata G, Saccone C, Serio G. A new method for calculating evolutionary substitution rates. *J Mol Evol.* 1984;20(1):86–93. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/6429346>.
- [18] Yang Z. Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: approximate methods. *J Mol Evol.* 1994 Sep;39(3):306–314.
- [19] Akaike H. A new look at the statistical model identification. *IEEE Trans Automat Contr.* 1974 Dec;19(6):716–723.
- [20] Schwarz G. Estimating the Dimension of a Model. *Ann Stat.* 1978 Mar;6(2):461–464.
- [21] Teng G, Papavasiliou FN. Immunoglobulin somatic hypermutation. *Annu Rev Genet.* 2007;41:107–120. Available from: <http://dx.doi.org/10.1146/annurev.genet.41.110306.130340>.
- [22] Aitchison J. Principal component analysis of compositional data. *Biometrika.* 1983 Apr;70(1):57–65. Available from: <http://dx.doi.org/10.1093/biomet/70.1.57>.
- [23] McCoy CO, Bedford T, Minin VN, Bradley P, Robins H, Matsen FA IV. figshare upload: Selection plots for all genes; 2015. Available from: <http://dx.doi.org/10.6084/m9.figshare.1399200>.
- [24] McCoy CO, Bedford T, Minin VN, Bradley P, Robins H, Matsen FA IV. figshare upload: Per-site per-gene estimates of selection posterior densities; 2015. Available from: <http://dx.doi.org/10.6084/m9.figshare.1399201>.
- [25] Lefranc MP, Pommié C, Ruiz M, Giudicelli V, Foulquier E, Truong L, et al. IMGT unique numbering for immunoglobulin and T cell receptor variable domains and Ig superfamily V-like domains. *Dev Comp Immunol.* 2003 Jan;27(1):55–77.
- [26] Leaver-Fay A, Tyka M, Lewis SM, Lange OF, Thompson J, Jacak R, et al. ROSETTA3: an object-oriented software suite for the simulation and design of macromolecules. *Methods Enzymol.* 2011;487:545–574.
- [27] Hollander M, Wolfe DA. *Nonparametric Statistical Methods.* Wiley; 1973.
- [28] Gerhardt S, Abbott WM, Hargreaves D, Paupit RA, Davies RA, Needham MRC, et al. Structure of IL-17A in complex with a potent, fully human neutralizing antibody. *J Mol Biol.* 2009 18 Dec;394(5):905–921.
- [29] Clark AG, Eisen MB, Smith DR, Bergman CM, Oliver B, Markow TA, et al. Evolution of genes and genomes on the Drosophila phylogeny. *Nature.* 2007;450(7167):203–218.
- [30] Lindblad-Toh K, Garber M, Zuk O, Lin MF, Parker BJ, Washietl S, et al. A high-resolution map of human evolutionary constraint using 29 mammals. *Nature.* 2011;478(7370):476–482.
- [31] Kosakovsky Pond SL, Scheffler K, Gravenor MB, Poon AFy, Frost SDw. Evolutionary Fingerprinting of Genes. *Mol Biol Evol.* 2010 Mar;27(3):520–536. Available from: <http://dx.doi.org/10.1093/molbev/msp260>.
- [32] Mirsky A, Kazandjian L, Anisimova M. Antibody-specific Model of Amino Acid Substitution for Immunological Inferences from Alignments of Antibody Sequences. *Mol Biol Evol.* 2014 21 Dec; Available from: <http://dx.doi.org/10.1093/molbev/msu340>.
- [33] Clarke SH, Huppi K, Ruezinsky D, Staudt L, Gerhard W, Weigert M. Inter- and intracolon diversity in the antibody response to influenza hemagglutinin. *J Exp Med.* 1985 1 Apr;161(4):687–704. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/3920342>.
- [34] Chang B, Casali P. The CDR1 sequences of a major proportion of human germline Ig VH genes are inherently susceptible to amino acid replacement. *Immunol Today.* 1994 Aug;15(8):367–373. Available from: [http://dx.doi.org/10.1016/0167-5699\(94\)90175-9](http://dx.doi.org/10.1016/0167-5699(94)90175-9).
- [35] Lossos IS, Tibshirani R, Narasimhan B, Levy R. The inference of antigen selection on Ig genes. *J Immunol.* 2000 1 Nov;165(9):5122–5126. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/11046043>.
- [36] Bose B, Sinha S. Problems in using statistical analysis of replacement and silent mutations in antibody genes for determining antigen-driven affinity selection. *Immunology.* 2005 Oct;116(2):172–183. Available from: <http://dx.doi.org/10.1111/j.1365-2567.2005.02208.x>.
- [37] Hershberg U, Uduman M, Shlomchik MJ, Kleinstein SH. Improved methods for detecting selection by mutation analysis of Ig V region sequences. *Int Immunol.* 2008 7 Apr;20(5):683–694. Available from: <http://dx.doi.org/10.1093/intimm/dxn026>.
- [38] Yaari G, Uduman M, Kleinstein SH. Quantifying selection in high-throughput Immunoglobulin sequencing data sets. *Nucleic Acids Res.* 2012 27 May;40(17):e134. Available from: <http://dx.doi.org/10.1093/nar/gks457>.

- [39] Steiman-Shimony A, Edelman H, Hutzler A, Barak M, Zuckerman NS, Shahaf G, et al. Lineage tree analysis of immunoglobulin variable-region gene mutations in autoimmune diseases: Chronic activation, normal selection. *Cell Immunol.* 2006 Dec;244(2):130–136. Available from: <http://dx.doi.org/10.1016/j.cellimm.2007.01.009>.
- [40] Abraham RS, Manske MK, Zuckerman NS, Sohni A, Edelman H, Shahaf G, et al. Novel analysis of clonal diversification in blood B cell and bone marrow plasma cell clones in immunoglobulin light chain amyloidosis. *J Clin Immunol.* 2006 28 Dec;27(1):69–87. Available from: <http://dx.doi.org/10.1007/s10875-006-9056-9>.
- [41] Barak M, Zuckerman N, Edelman H, Unger R, Mehr R. IgTree (c) : Creating Immunoglobulin variable region gene lineage trees. *Journal of Immunological Methods.* 2008;338(1-2):67–74. Available from: <http://dx.doi.org/10.1016/j.jim.2008.06.006>.
- [42] Shahaf G, Barak M, Zuckerman NS, Swerdlin N, Gorfine M, Mehr R. Antigen-driven selection in germinal centers as reflected by the shape characteristics of immunoglobulin gene lineage trees: a large-scale simulation study. *J Theor Biol.* 2008 13 Aug;255(2):210–222. Available from: <http://dx.doi.org/10.1016/j.jtbi.2008.08.005>.
- [43] Liberman G, Benichou J, Tsaban L, Glanville J, Louzoun Y. Multi Step Selection in Ig H Chains is Initially Focused on CDR3 and Then on Other CDR Regions. *Front Immunol.* 2013 17 Sep;4:274. Available from: <http://dx.doi.org/10.3389/fimmu.2013.00274>.
- [44] Uduman M, Shlomchik MJ, Vigneault F, Church GM, Kleinstein SH. Integrating B Cell Lineage Information into Statistical Tests for Detecting Selection in Ig Sequences. *J Immunol.* 2014 1 Feb;192(3):867–874. Available from: <http://dx.doi.org/10.4049/jimmunol.1301551>.
- [45] Mooers AO, Heard SB. Inferring Evolutionary Process from Phylogenetic Tree Shape. *Q Rev Biol.* 1997 Mar;72(1):31–54. Available from: <http://dx.doi.org/10.2307/3036810>.
- [46] Drummond AJ, Suchard MA. Fully Bayesian tests of neutrality using genealogical summary statistics. *BMC Genet.* 2008 31 Oct;9:68. Available from: <http://dx.doi.org/10.1186/1471-2156-9-68>.
- [47] Li H, Wiehe T. Coalescent tree imbalance and a simple test for selective sweeps based on microsatellite variation. *PLoS Comput Biol.* 2013 16 May;9(5):e1003060. Available from: <http://dx.doi.org/10.1371/journal.pcbi.1003060>.
- [48] Łuksza M, Lässig M. A predictive fitness model for influenza. *Nature.* 2014 26 Feb; Available from: <http://dx.doi.org/10.1038/nature13087>.
- [49] Neher RA, Russell CA, Shraiman BI. Predicting evolution from the shape of genealogical trees. *Elife.* 2014 11 Nov;3. Available from: <http://dx.doi.org/10.7554/eLife.03568>.
- [50] Ralph DK, Matsen FA IV. Consistency of VDJ rearrangement and substitution parameters enables accurate B cell receptor sequence annotation. *arXiv; Submitted to PLOS Computational Biology.* 2015 13 Mar; Available from: <http://arxiv.org/abs/1503.04224>.
- [51] Liao HX, Lynch R, Zhou T, Gao F, Alam SM, Boyd SD, et al. Co-evolution of a broadly neutralizing HIV-1 antibody and founder virus. *Nature.* 2013 25 Apr;496(7446):469–476. Available from: <http://dx.doi.org/10.1038/nature12053>.
- [52] Bashford-Rogers RJM, Palser AL, Huntly BJ, Rance R, Vassiliou GS, Follows GA, et al. Network properties derived from deep sequencing of human B-cell receptor repertoires delineate B-cell populations. *Genome Res.* 2013 6 Jun;23(11):1874–1884. Available from: <http://dx.doi.org/10.1101/gr.154815.113>.
- [53] Carlson CS, Emerson RO, Sherwood AM, Desmarais C, Chung MW, Parsons JM, et al. Using synthetic templates to design an unbiased multiplex PCR assay. *Nat Commun.* 2013;4:2680.
- [54] Robins HS, Campregher PV, Srivastava SK, Wachter A, Turtle CJ, Khsai O, et al. Comprehensive assessment of T-cell receptor  $\beta$ -chain diversity in  $\alpha\beta$  T cells. *Blood.* 2009 25 Aug;114(19):4099–4107. Available from: <http://dx.doi.org/10.1182/blood-2009-04-217604>.
- [55] Gotoh O. An improved algorithm for matching biological sequences. *J Mol Biol.* 1982 15 Dec;162(3):705–708.
- [56] Xu JL, Davis MM. Diversity in the CDR3 Region of VH Is Sufficient for Most Antibody Specificities. *Immunity.* 2000 1 Jul;13(1):37–45.
- [57] Zhu J, Ofek G, Yang Y, Zhang B, Louder MK, Lu G, et al. Mining the antibodyome for HIV-1-neutralizing antibodies with next-generation sequencing and phylogenetic pairing of heavy/light chains. *Proceedings of the National Academy of Sciences.* 2013 Mar;.
- [58] Jiang N, He J, Weinstein JA, Penland L, Sasaki S, He XS, et al. Lineage structure of the human antibody repertoire in response to influenza vaccination. *Sci Transl Med.* 2013 6 Feb;5(171):171ra19. Available from: <http://dx.doi.org/10.1126/scitranslmed.3004794>.

- [59] Kepler TB. Reconstructing a B-cell clonal lineage. I. Statistical inference of unobserved ancestors. *F1000Res*. 2013 3 Apr;2:103. Available from: <http://dx.doi.org/10.12688/f1000research.2-103.v1>.
- [60] Felsenstein J. Evolutionary trees from DNA sequences: a maximum likelihood approach. *J Mol Evol*. 1981;17(6):368–376. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/7288891>.
- [61] Tavaré S. Some probabilistic and statistical problems in the analysis of DNA sequences. *Lectures on mathematics in the life sciences*. 1986;.
- [62] Yang Z. Maximum likelihood phylogenetic estimation from DNA sequences with variable rates over sites: approximate methods. *J Mol Evol*. 1994 Sep;39(3):306–314. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/7932792>.
- [63] Salemi M, Lemey P, Vandamme AM. *The Phylogenetic Handbook: A Practical Approach to Phylogenetic Analysis and Hypothesis Testing*. Cambridge University Press; 2009.
- [64] Felsenstein J. *Inferring phylogenies*. Sinauer Associates; 2004.
- [65] Guéguen L, Gaillard S, Boussau B, Gouy M, Groussin M, Rochette NC, et al. Bio++: efficient extensible libraries and tools for computational molecular evolution. *Mol Biol Evol*. 2013 21 May;30(8):1745–1750.
- [66] Ayres DL, Darling A, Zwickl DJ, Beerli P, Holder MT, Lewis PO, et al. BEAGLE: an application programming interface and high-performance computing library for statistical phylogenetics. *Syst Biol*. 2011 1 Oct;61(1):170–173.
- [67] Powell M. *The BOBYQA algorithm for bound constrained optimization without derivatives*. Cambridge NA Report NA2009/06, University of Cambridge, Cambridge. 2009;.
- [68] Johnson SG. *The NLOpt nonlinear-optimization package*; 2010.
- [69] Brent RP. *Algorithms for Minimization Without Derivatives*. Courier Dover Publications; 1973.
- [70] van den Boogaart KG, Tolosana-Delgado R. “Compositions”: a unified R package to analyze compositional data. *Computers & Geosciences*. 2008;34(4):320–338.
- [71] Goldman N, Yang Z. A codon-based model of nucleotide substitution for protein-coding DNA sequences. *Mol Biol Evol*. 1994 Sep;11(5):725–736. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/7968486>.
- [72] Muse SV, Gaut BS. A likelihood approach for comparing synonymous and nonsynonymous nucleotide substitution rates, with application to the chloroplast genome. *Mol Biol Evol*. 1994 Sep;11(5):715–724. Available from: <http://www.ncbi.nlm.nih.gov/pubmed/7968485>.
- [73] Murrell B, Moola S, Mabona A, Weighill T, Sheward D, Kosakovsky Pond SL, et al. FUBAR: A Fast, Unconstrained Bayesian AppRoximation for Inferring Selection. *Mol Biol Evol*. 2013 May;30(5):1196–1205. Available from: <http://dx.doi.org/10.1093/molbev/mst030>.
- [74] Robbins H. *An Empirical Bayes Approach to Statistics*. In: *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*. The Regents of the University of California; 1956. .
- [75] Hasegawa M, Kishino H, Yano T. Dating of the human-ape splitting by a molecular clock of mitochondrial DNA. *J Mol Evol*. 1985;22(2):160–174.
- [76] Bielejec F, Lemey P, Carvalho LM, Baele G, Rambaut A, Suchard MA. piBUSS: a parallel BEAST/BEAGLE utility for sequence simulation under complex evolutionary scenarios; 2013.
- [77] Drummond AJ, Suchard MA, Xie D, Rambaut A. Bayesian Phylogenetics with BEAUti and the BEAST 1.7. *Mol Biol Evol*. 2012 Aug;29(8):1969–1973. Available from: <http://dx.doi.org/10.1093/molbev/mss075>.
- [78] Zaharia M, Chowdhury M, Franklin M, Shenker S, Stoica I. Spark: cluster computing with working sets. In: Nahum E, Xu D, editors. *2nd USENIX conference on hot topics in cloud computing*. static.usenix.org; 2010. p. 10. Available from: [http://static.usenix.org/legacy/events/hotcloud10/tech/full\\_papers/Zaharia.pdf](http://static.usenix.org/legacy/events/hotcloud10/tech/full_papers/Zaharia.pdf).
- [79] Berman HM, Westbrook J, Feng Z, Gilliland G, Bhat TN, Weissig H, et al. The Protein Data Bank. *Nucleic Acids Res*. 2000 1 Jan;28(1):235–242.
- [80] Altschul SF, Madden TL, Schäffer AA, Zhang J, Zhang Z, Miller W, et al. Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic Acids Res*. 1997 1 Sep;25(17):3389–3402.
- [81] Delano WL. *The PyMOL Molecular Graphics System*. citeulike.org; 2002.
- [82] Holm S. A simple sequentially rejective multiple test procedure. *Scand Stat Theory Appl*. 1979;.
- [83] Amazon Web Services. *Amazon EC2 FAQs*; 2014. Available from: [http://aws.amazon.com/ec2/faqs/#What\\_is\\_Amazon\\_Elastic\\_Compute\\_Cloud\\_Amazon\\_EC2](http://aws.amazon.com/ec2/faqs/#What_is_Amazon_Elastic_Compute_Cloud_Amazon_EC2).



## Supplementary Material: Quantifying evolutionary constraints on B cell affinity maturation

CONNOR O. MCCOY, TREVOR BEDFORD, VLADIMIR N. MININ, PHILIP  
BRADLEY, HARLAN ROBINS, AND FREDERICK A. MATSEN IV  
*Philosophical Transactions of the Royal Society B*; doi:10.1098/rstb.2014-0244

### SUPPLEMENTARY METHODS

**Derivation of the Gamma-Poisson marginal likelihood with varying observation depth.** We will use the same notation as in the Materials and Methods section. Our first task is to write down a likelihood of  $\alpha$  and  $\beta$  given a collection of counts. To do so we will marginalize out the rates  $\lambda_l$  when they are drawn from a  $\text{Gamma}(\alpha, \beta)$  as in the main text.

The likelihood for a single site is (omitting  $l$  for now):

$$\begin{aligned} P(C|t, \alpha, \beta) &= \int_0^\infty P(C|t, \lambda) P(\lambda|\alpha, \beta) d\lambda \\ &= \int_0^\infty \frac{(\lambda t)^C e^{-\lambda t}}{C!} P(\lambda|\alpha, \beta) d\lambda \\ &= \int_0^\infty \frac{(\lambda t)^C e^{-\lambda t}}{C!} \left[ \frac{\beta^\alpha}{\Gamma(\alpha)} \lambda^{\alpha-1} e^{-\beta\lambda} \right] d\lambda \\ &= \frac{\beta^\alpha t^C}{C! \Gamma(\alpha)} \int_0^\infty \lambda^{C+\alpha-1} e^{-\lambda(t+\beta)} d\lambda. \end{aligned}$$

Letting  $\alpha' = C + \alpha$  and  $\beta' = t + \beta$ , introduce a normalizing constant for the distribution  $\text{Gamma}(\alpha', \beta')$ :

$$\begin{aligned} P(C|t, \alpha, \beta) &= \frac{\beta^\alpha t^C}{C! \Gamma(\alpha)} \frac{\Gamma(\alpha')}{\beta'^{\alpha'}} \int_0^\infty \frac{\beta'^{\alpha'}}{\Gamma(\alpha')} \lambda^{\alpha'-1} e^{-\lambda\beta'} d\lambda \\ &= \frac{\beta^\alpha t^C}{C! \Gamma(\alpha)} \frac{\Gamma(\alpha')}{\beta'^{\alpha'}} \int_0^\infty \text{DGamma}(\lambda; \alpha', \beta') d\lambda. \end{aligned}$$

The integral over the support of the Gamma distribution is 1, so:

$$\begin{aligned} P(C|t, \alpha, \beta) &= \frac{\beta^\alpha t^C}{C! \Gamma(\alpha)} \frac{\Gamma(\alpha')}{\beta'^{\alpha'}} \\ &= \frac{\beta^\alpha t^C}{C! \Gamma(\alpha)} \frac{\Gamma(C + \alpha)}{(t + \beta)^{C+\alpha}}. \end{aligned}$$

The overall marginal likelihood is the product over such sites:

$$\begin{aligned}
\mathcal{L} &= P(C_1, \dots, C_L | t_1, \dots, t_L, \alpha, \beta) = \prod_l \frac{\beta^\alpha t_l^{C_l}}{C_l! \Gamma(\alpha)} \frac{\Gamma(C_l + \alpha)}{(t_l + \beta)^{C_l + \alpha}} \\
&= \left( \frac{\beta^\alpha}{\Gamma(\alpha)} \right)^L \prod_l \frac{t_l^{C_l}}{C_l!} \frac{\Gamma(C_l + \alpha)}{(t_l + \beta)^{C_l + \alpha}} \\
&= \left( \frac{\beta^\alpha}{\Gamma(\alpha)} \right)^L \prod_l \frac{t_l^{C_l}}{\Gamma(C_l + 1)} \frac{\Gamma(C_l + \alpha)}{(t_l + \beta)^{C_l + \alpha}},
\end{aligned}$$

giving (1).

*Posterior for  $\lambda$ .* Our eventual goal is a regularized posterior estimate of the rates  $\lambda_l$ . For a single site, once again dropping  $l$ :

$$P(\lambda | C, t, \hat{\alpha}, \hat{\beta}) \propto P(C | \lambda, t) P(\lambda | \hat{\alpha}, \hat{\beta}).$$

Substituting in the PDFs for the distributions employed for  $C$  and  $\lambda$ :

$$P(\lambda | C, t, \hat{\alpha}, \hat{\beta}) \propto \frac{\hat{\beta}^{\hat{\alpha}} t^C}{C! \Gamma(\hat{\alpha})} \lambda^{C + \hat{\alpha} - 1} e^{-\lambda(t + \hat{\beta})}.$$

As in the main text, we let  $\hat{\alpha}' = C + \hat{\alpha}$  and  $\hat{\beta}' = t + \hat{\beta}$ .

$$\begin{aligned}
P(\lambda | C, t, \hat{\alpha}, \hat{\beta}) &\propto \frac{\hat{\beta}^{\hat{\alpha}} t^C}{C! \Gamma(\hat{\alpha})} \frac{\Gamma(\hat{\alpha}')}{\hat{\beta}'^{\hat{\alpha}'}} \left[ \frac{\hat{\beta}'^{\hat{\alpha}'}}{\Gamma(\hat{\alpha}')} \lambda^{\hat{\alpha}' - 1} e^{-\lambda(\hat{\beta}')} \right] \\
&\propto \frac{\hat{\beta}^{\hat{\alpha}} t^C}{C! \Gamma(\hat{\alpha})} \frac{\Gamma(\hat{\alpha}')}{\hat{\beta}'^{\hat{\alpha}'}} \text{DGamma}(\lambda; \hat{\alpha}', \hat{\beta}') \\
&\propto \text{DGamma}(\lambda; \hat{\alpha}', \hat{\beta}'),
\end{aligned}$$

hence these two probability densities are equal, justifying (2).

individual	raw	sequence count	
		unique by well	unique overall
A	52,381,123	8,275,848	4,778,427
B	59,241,547	9,820,657	5,826,068
C	66,469,248	8,452,997	4,419,453

TABLE S1. Number of memory BCR sequences obtained by individual. “raw” refers to the number of reads obtained from sequencing, “unique by well” the number of unique sequences after performing clustering on reads for each barcoded PCR well, and “unique overall” the total number of unique sequences in the sample.

individual	cell type	in-frame	out-of-frame
A	naïve	0.93	0.92
	memory	0.08	0.15
B	naïve	0.91	0.89
	memory	0.22	0.27
C	naïve	0.92	0.90
	memory	0.20	0.28

TABLE S2. Fraction of BCR sequences that were identical to germline in the regions inferred to derive from germline. The fractions are stratified by individual, cell type, and frame status.

		Individual A				Individual B				Individual C			
		A	G	C	T	A	G	C	T	A	G	C	T
IGHV	germline	0.283	0.27	0.255	0.192	0.279	0.27	0.261	0.19	0.285	0.268	0.258	0.189
	sequence	0.277	0.261	0.256	0.206	0.276	0.266	0.261	0.197	0.282	0.265	0.258	0.196
IGHD	germline	0.199	0.328	0.141	0.332	0.196	0.323	0.157	0.324	0.197	0.326	0.153	0.324
	sequence	0.197	0.315	0.168	0.321	0.198	0.309	0.176	0.317	0.197	0.314	0.172	0.317
IGHJ	germline	0.197	0.428	0.22	0.154	0.2	0.424	0.223	0.154	0.186	0.438	0.225	0.151
	sequence	0.186	0.433	0.222	0.159	0.193	0.427	0.224	0.156	0.18	0.44	0.227	0.153

TABLE S3. Empirical stationary distribution for germline and observed sequences.

#### SUPPLEMENTARY FIGURES AND TABLES




FIGURE S1. GTR coefficients for the  $t_r Q_i \Gamma_i$  model estimated under maximum likelihood. Rows index the nucleotide found in the germline sequence, whereas columns index the nucleotide found in the observed sequence.




FIGURE S2. Pairwise comparison of off-diagonal entries in maximum-likelihood  $Q$  matrices under the  $t_r Q_i \Gamma_i$  model between the three individuals. Coefficients are shown in Fig. S1.




FIGURE S3. Comparison of Hamming distance between V genes (x-axis) and Euclidean distance between centered log-transformed median time transition matrices for productive rearrangements (y-axis). Colors indicate whether the V genes in a comparison come from the same or different subgroups. The correlation between the two was significant ( $p < 10^{-15}$ , Spearman's  $\rho = 0.197$ ).



FIGURE S4. Amino acid profiles of out-of-frame and functional B cell sequences as aligned by the IMGT alignment. Top panel: frequency of amino acids per site. Letters to the left of the line show the profile for out-of-frame sequences and those to the right of the line show the profile for functional sequences. Bottom panel: amino acid frequency in functional sequences divided by that in out-of-frame sequences.




FIGURE S5. Pairwise comparisons of site-specific  $\omega$  estimates between the three individuals along with the  $R^2$  value from a linear model fit using  $\log_{10}(\omega)$  for both the predictor (x-axis) and response (y-axis).




FIGURE S6. Top panel: site-specific  $\omega$  estimates under simulated data with varying coverage. Inverted triangles show sites where the germline state was Tryptophan or Methionine, from which no synonymous changes are possible. Dashed black line shows simulated  $\omega$ . Middle panel: proportion (second panel) of mutations at each position which were nonsynonymous (N) or synonymous (S). Bottom panel: sequence coverage by codon position.




FIGURE S7. Normalized solvent-accessible surface area (SASA) values by per-site  $\omega$  classification. A SASA value of 1.0 indicates that the residue is fully exposed, while a value of 0.0 indicates that the residue is buried. Sites under negative selection are significantly less exposed than sites under positive selection ( $p < 10^{-12}$ ) or neutral selection ( $p < 10^{-15}$ ) by Bonferroni-corrected Wilcoxon rank-sum test. Neutral sites were less exposed than sites under positive selection ( $p < 0.002$ ).